Noname manuscript No.
(will be inserted by the editor)

Sensor Selection by Linear Programming

Joseph Wang · Kirill Trapeznkov · Venkatesh Saligrama

arXiv:1509.02954v1 [stat.ML] 9 Sep 2015

Received: date / Accepted: date

Abstract We learn sensor trees from training data to minimize sensor acquisition costs during test time. Our system adaptively selects sensors at each stage if necessary to make a confident classification. We pose the problem as empirical risk minimization over the choice of trees and node decision rules. We decompose the problem, which is known to be intractable, into combinatorial (tree structures) and continuous parts (node decision rules) and propose to solve them separately. Using training data we greedily solve for the combinatorial tree structures and for the continuous part, which is a non-convex multilinear objective function, we derive convex surrogate loss functions that are piecewise linear. The resulting problem can be cast as a linear program and has the advantage of guaranteed convergence, global optimality, repeatability and computational efficiency. We show that our proposed approach outperforms the state-of-art on a number of benchmark datasets.

Keywords Adaptive sensor selection · Resource-constrained learning · Test-time budgeted learning

1 Introduction

Many scenarios involve classification systems constrained by measurement acquisition budget. In this setting, a collection of sensor modalities with varying costs are available to the decision system. Our goal is to learn adaptive decision rules from labeled training data that when presented with a new unseen example would

Joseph Wang
8 St. Mary's Street, Boston, MA 02215, USA
Tel.: (617) 353-2811
Fax: (617) 353-7337
E-mail: joewang@bu.edu

Kirill Trapeznikov
E-mail: kirill.trapeznikov@stresearch.com

Venkatesh Saligrama
E-mail: srv@bu.edu

select the most informative and cost-effective strategy for the example. In contrast to non-adaptive methods Efron et al (2004); Xu et al (2012), which attempt to identify a common sparse subset of sensors that can work well for all examples, our goal is an adaptive method that can classify typical cases using inexpensive sensors and using expensive sensors only for atypical cases.

We propose to learn a sensor tree using labeled training examples for making decisions on unseen test examples. The learned sensor tree is composed of internal node decision rules. Given an example these decision rules select sensors and guide an example along a particular path terminating at a leaf where it is classified with a classifier. We pose the problem as a global empirical risk minimization (ERM) over the choice of tree structures, node decision rules and leaf classifiers.

The general problem is a highly coupled problem consisting of combinatorial (sensor tree structure) and continuous components (decision rules to generalize to unseen examples) and difficult to optimize. To gain further insight we abstract away the generalization aspect and observe that the resulting combinatorial problem, a special case of ours, is known to be NP hard and requires greedy approximations Chakaravarthy et al (2007); Cicalese et al (2014).

The combinatorial issue can be circumvented in cases where expert knowledge exists, only a few sensors (as in Trapeznikov et al (2013); Trapeznikov and Saligrama (2013); Wang et al (2014b)) exist or for small depth trees. For these latter two cases we contruct an exhaustive tree and globally learn decision functions using a linear program by generalizing the cascade structures presented in Trapeznikov and Saligrama (2013); Wang et al (2014b) to binary trees, resulting in more flexible decision systems. Convex surrogates of products of indicators have previously been studied for supervised learning Wang and Saligrama (2013).

For more general cases we propose a two-step approach to decouple the issue of sensor structure from the decision rule design.

- We greedily solve for the combinatorial tree structure and obtain feature/sensor sub-collections efficiently by a greedy approximation to the NP-hard problem. From these subsets, we construct a binary tree using hierarchical clustering of feature subsets.
- On the learned tree structure, our problem now reduces to the ERM problem discussed above for a fixed tree where we apply a novel surrogate, allowing us to jointly learn the decision functions in the tree by solving a linear program.

In the experiments, we demonstrate performance of our approach both for when feature subsets and tree structure is given and when feature subsets and tree structure must be learned. We show on real world data that our approach outperforms previously proposed approaches to budgeted learning.

1.1 Related Work

There is an extensive literature on adaptive methods for sensor selection for reducing test-time costs. It arguably originated with detection cascade (see Zhang and Zhang (2010); Chen et al (2012) and references therein), which is a popular method in reducing computation cost in object detection for cases with highly skewed class imbalance and generic features. Computationally cheap features are used at first

to filter out negative examples and the more expensive features are used in the later stages.

Our technical approach is closely related to Trapeznikov and Saligrama (2013) and Wang et al (2014b). Like us they formulate an ERM problem and generalize detection cascades to classifier cascades and handle balanced and/or multi-class scenarios. Like us, Wang et al (2014b) construct convex surrogates for their empirical risk functions and propose efficient LP solutions. Unlike us their approach is limited to cascades of known structure and cannot handle trees and unknown sensor structures.

Conceptually, our work is closely related to Xu et al (2013) and Kusner et al (2014), who introduced cost-sensitive tree of classifiers (CSTC) for reducing test time costs. Like our paper they proposed a global ERM problem. They solve for the tree structure, internal decision rules and leaf classifiers jointly using alternative minimization techniques. Recently, Kusner et al (2014) propose a more efficient version of CSTC. In contrast we decompose our global objective and separately solve the individual parts. The disadvantage of our decoupled approach is somewhat offset by globally convergent solution for the decision rules once a structure is determined. Nevertheless, which approach is better is an important question that must be addressed but outside the scope of this work.

The subject of this paper is broadly related to other adaptive methods in the literature but unlike us these methods do not learn sensor trees but learn policies from training data. Generative methods pose the problem as a POMDP, learn conditional probability models Zubek and Dietterich (2002); Sheng and Ling (2006); Bilgic and Getoor (2007); Ji and Carin (2007); Kanani and Melville (2008); Kapoor and Horvitz (2009); Gao and Koller (2011) and myopically select features based information gain of unknown features. MDP-based methods Karayev et al (2013), Dulac-Arnold et al (2011), He et al (2012), Busa-Fekete et al (2012) encode current observations as state, unused features as action space, and formulate various reward functions to account for classification error and costs. He et. al. He et al (2012) apply imitation learning of a greedy policy with a single classification step as actions. Dulac-Arnold et al (2011) and Karayev et al (2013) apply reinforcement learning to solve this MDP. Busa-Fekete et al (2012) propose classifier cascades within an MDP framework. They consider a fixed-ordering of features and extend sequential boosted classifier with an additional skip action.

2 Problem Formulation: Global Empirical Risk Minimization Objective

We consider learning an adaptive decision system with training examples (x_1, y_1), $\ldots, (x_N, y_N)$ with L sensors and acquisition cost c_m, $m = 1, 2, \ldots, L$. We can pose the problem of learning a rooted sensor tree as an ERM problem. Our system is composed of three components, a tree, T, decision rules $\mathbf{g}_J = \{g_j\}_{j=1}^{J}$, $\mathbf{g}_j \in \mathcal{G}_j$ associated with J internal nodes, and classifier $\mathbf{f}_K = \{f_k\}_{k=1}^{K}$, $f_k \in \mathcal{F}_k$ associated with leaves. Each internal node of the tree is associated with a sensor and its children denote available sensor choices. Each leaf, k, is associated with a sensor subset, S_k, and corresponds to the unique path from the root to the leaf. The decision rule, g_j, associated with node j acts upon acquired measurements for an example and routes it to one of its children. By uniquely associating each leaf, k,

with the sensor outputs b_k we can write ERM as:

$$L(\mathcal{S}, \mathbf{f}, \mathbf{g}) = \sum_{i=1}^{N} \sum_{k=1}^{K} \left(\mathbb{1}_{f_k(x_i) \neq y_i} + \alpha \sum_{m \in S_k} c_m \right) \mathbb{1}_{\mathbf{g}_J(x_i) = b_k} \\ \longrightarrow \min_{\mathcal{S}} \min_{\mathbf{f}_K} \min_{\mathbf{g}_J} L(\mathcal{S}, \mathbf{f}, \mathbf{g}) \tag{1}$$

where α is a trade-off parameter balancing classification performance with sensor acquisition cost and $\mathcal{S}$ is the set of paths. An instance of this general problem has been considered in the literature and shown to be NP hard Chakaravarthy et al (2007) for the special case of discrete valued sensor measurements, arbitrarily powerful decision rules, $\mathbf{g}_J$ and with separable leaves, i.e, features acquired corresponding to the leaf path uniquely & correctly identifies the class. The authors develop greedy algorithms for approximating their solution. In our setting we allow for continous valued high-dimensional measurements and so we cannot discretize our space. Furthermore we are not in a separable situation and the Bayes error is not zero. However, our general ERM problem is highly coupled, difficult to optimize. This motivates our proposed decomposition approach described below.

Indeed, assuming arbitrary families $\mathcal{G}_j$ in Eq. 1 leads to new insights that motivates our decomposition approach. The general problem reduced to a purely combinatorial structure learning with arbitrary $\mathcal{G}_j$ when we also have access to an oracle classifier $\mathbf{f}_K$ capable of classifying with any subset of acquired features. The resulting problem while NP hard as before is amenable to greedy strategies. Alternatively, given a tree structure and the oracle classifier the optimization objective takes a multilinear form as evidenced by Eq. 1 which also lends itself to optimization strategies. The only issue remaining is that of an oracle classifier, which can be circumvented for our approach. This overall perspective justifies our approach:

1. Learn the tree structure assuming powerful decision rules.
2. Learn decision functions $g_1, \ldots, g_{K-1}$ for the tree structure learned in Step 1.

2.1 Learning Tree Structures Greedily

For simplicity we assume a binary sensor tree with K leaves and $K-1$ internal nodes. Motivated by previous methods Gao and Koller (2011); Xu et al (2013) we assume that the number of leaves K is small relative to the feature dimension. Our approach identifies a sub-collection of subsets of features $\mathcal{S} = \{S_1, \ldots, S_K\}$ from training data, and as such, the tree structure, T. Assuming arbitrarily powerful decision functions, $g_1, \ldots, g_{K-1}$, effectively implies we can route each example to its optimal subset. Furthermore, assuming that we have access to an oracle classifer, f_j's we can predicted the class on any subset of features. Then the optimization loss of Eq. 1 associated with the subcollection $\mathcal{S}$ reduces to:

$$L(S_1, \ldots, S_K) = \frac{1}{N} \sum_{i=1}^{N} \min_{j \in \{1, \ldots, K\}} \left[\mathbb{1}_{f_{s_j}(x_i) \neq y_i} + \alpha \sum_{k \in S_j} c_k \right] \tag{2}$$

Even in the absence of noise, as noted before the problem of minimizing this loss is NP-hard and motivates greedy strategies. Additionally, we overcome the

issue of an oracle classifier by learning it as we grow the tree greedily. While many greedy strategies exist in the literature for related objectives, they are not directly applicable to our setting[1]. Our greedy algorithm is related to Awasthi et al (2013), who learn sparse trees/polynomials in the separable PAC setting and provide statistical and computational guarantees. We adapt their approach to our setting as follows: given a sub-collection of features we expand this subcollection if there is a sensor that can further reduce our loss. If no sensor is found we restart the process and look for a new subset of sensors.

Algorithm 1 Sensor Subset Selection

Input: Number of Subsets K
Output: Feature subsets, $s_1, \ldots, s_K$
Initialize: $s_1, \ldots, s_K = \emptyset$
for k=1,. . . ,K **do**
  $j = \text{argmin}_{j \in s_k^C} \mathcal{L}(s_1, ..., s_k \cup j, ..., s_K)$
  while $L(s_1, ..., s_k \cup j, ..., s_K) < \mathcal{L}(s_1, ..., s_k, ..., s_K)$ **do**
    $s_k = s_k \cup j$
    $j = \text{argmin}_{j \in s_k^C} \mathcal{L}(s_1, ..., s_k \cup j, ..., s_K)$
  end while
end for

Tree Structure: Given the set of sensor subsets $s_1, \ldots, s_K$, the problem of choosing a tree structure partitioning between these sensor subsets arises. We propose a hierarchical clustering approach, where subsets are grouped based on the number of common elements. Given a set of feature subsets, the two subsets with the highest number of common elements are grouped together and replaced in the set with the intersection of their elements. This is recursively repeated until only a single subset exists, resulting in a binary tree structure. This can be viewed as a generalization of the cascade approach, as given the set of feature subsets where an additional sensor is added to each previous subset, a cascade structure is always recovered.

Once a tree structure, T is learned we need to populate it with decision functions so that we can generalize to unseen examples. Note that the learned sensor structure provides us with possible choices but does not tell us what choice to make on an unseen example. This motivates the following section.

2.2 Empirical Risk Problem for a Fixed Tree

We represent our decision system as a binary tree. The binary tree is composed of K leafs and $K-1$ internal nodes. At each internal node, $j = 1, \ldots, K-1$, is a binary decision function, $\text{sign}[g_j(\mathbf{x})] \in \{+1, -1\}$. This function determines which action should be taken for a given example. The binary decisions, $g_j(\mathbf{x})$'s, represent actions from the following set: stop and classify with the current set of measurements, or choose which sensor to acquire next. Each leaf node, $k = 1, \ldots K$, represents a terminal decision to stop and classify based on the available

[1] For instance Cicalese et al (2014) proposes a submodular surrogate to leverage properties of Wolsey greedy algorithm but their surrogates require discrete sensor measurements.

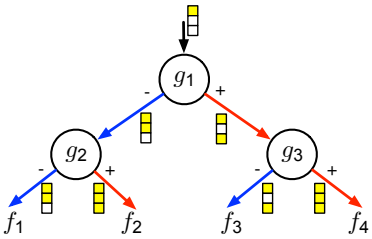

$$\underset{\mathbf{P}}{\begin{bmatrix}0&0&0\\0&1&0\\1&0&0\\1&0&1\end{bmatrix}} - \underset{\mathbf{N}}{\begin{bmatrix}1&1&0\\1&0&0\\0&0&1\\0&0&0\end{bmatrix}} = \begin{bmatrix}-1&-1&0\\-1&+1&0\\+1&0&-1\\+1&0&+1\end{bmatrix}$$

Fig. 1 An example decision system of depth two: node $g_1(x_1)$ selects either to acquire sensor 2 for a cost c_2 or 3 for a cost c_3. Node $g_2(x_1, x_2)$ selects either to stop and classify with sensors $\{1, 2\}$ or to acquire 3 for c_3 and then stop. Node $g_3(x_1, x_3)$ selects to classify with $\{1, 3\}$ or with $\{1, 2, 3\}$.

information[2] We assume that the leaf classifiers, f_j are known and fixed[3]. However note, the functions implicitly operate only on the sensors that have been acquired along the path to the corresponding node. The objective is to learn the internal decision functions: $g_j(\mathbf{x})$'s. We define the system risk:

$$R(\mathbf{g}, \mathbf{x}, y) = \sum_{k=1}^{K} R_k(f_k, \mathbf{x}, y) G_k(\mathbf{g}, \mathbf{x}) \tag{3}$$

Here, $\mathbf{g} = \{g_1, \dots g_{K-1}\}$ is the set of decision functions. $R_k(f_k, \mathbf{x}, y) = \mathbb{1}_{f_k(\mathbf{x}) \neq y} + \alpha \sum_{m \in S_k} c_m$ is the risk of making a decision at a leaf k. It consists of two terms: error of the classifier at the leaf and the cost of sensors acquired along the path from the root node to the leaf. S_k is this set of sensors, and α is a parameter that controls trade-off between acquisition cost and classification error. $G_k(\mathbf{g}, \mathbf{x}) \in \{0, 1\}$ is a binary state variable indicating if $\mathbf{x}$ is classified at the kth leaf. We compactly encode the path from the root to every leaf in terms of internal decisions, $g_j(\mathbf{x})$'s, by two auxiliary binary matrices: $\mathbf{P}$, $\mathbf{N} \in \{0, 1\}^{K \times K-1}$. If $\mathbf{P}_{k,j} = 1$ then, on the path to leaf k, a decision node j must be positive: $g_j > 0$. If $\mathbf{N}_{k,j} = 1$ then on the path to leaf k, a decision at node j must be negative: $g_j \leq 0$. A kth row in $\mathbf{P}$ and $\mathbf{N}$ jointly encode a path from the root node to a leaf k. The sign pattern for each path is obtained by $\mathbf{P} - \mathbf{N}$. For an example refer to Fig. 1. Using this path matrix, the state variable can be defined: $G_k(\mathbf{g}, \mathbf{x}) = \prod_{j=1}^{K-1} [\mathbb{1}_{g_j(\mathbf{x})>0}]^{\mathbf{P}_{k,j}} [\mathbb{1}_{g_j(\mathbf{x})\leq 0}]^{\mathbf{N}_{k,j}}$.

Our goal is to learn decision functions $g_1, \dots, g_{K-1}$ that minimize the expected system risk: $\min_{\mathbf{g}} \mathbb{E}_{\mathcal{D}}[R(\mathbf{g}, \mathbf{x}, y)]$ However, the model $\mathcal{D}$ is assumed to be unknown and cannot be estimated reliably due to potentially high-dimensional nature of sensor outputs. Instead, we are given a set of N training examples with full sensor measurements, $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_N, y_N)$. We approximate the expected risk by a sample average over the data and construct the following ERM problem:

$$\min_{\mathbf{g}} \sum_{i=1}^{N} R(\mathbf{g}, \mathbf{x}_i, y_i) = \sum_{i=1}^{N} \sum_{k=1}^{K} \overbrace{R_k(f_k, \mathbf{x}_i, y_i)}^{\text{risk of leaf } k} \underbrace{\prod_{j=1}^{K-1} [\mathbb{1}_{g_j(\mathbf{x}_i)>0}]^{\mathbf{P}_{k,j}} [\mathbb{1}_{g_j(\mathbf{x}_i)\leq 0}]^{\mathbf{N}_{k,j}}}_{G_k(\cdot) \,=\, \text{state of } \mathbf{x}_i \text{ in a tree}} \tag{4}$$

[2] For notational simplicity, we denote applying a decision node and a leaf classifier as $g_j(\mathbf{x})$ and $f_k(\mathbf{x})$ respectively.

[3] Note that the classifiers at each leaf, $f_k(\mathbf{x}) \in \mathcal{Y}$, can be learned for the K leaves once the tree structure, T, has been determined from the previous section.

Note that by the definition of risk in (3), the ERM problem can be viewed as a minimization over a function of indicators with respect to decisions: $g_1(\mathbf{x}), \ldots, g_{K-1}(\mathbf{x})$.

3 Convex Re-formulation and Solution by Linear Programming

A popular approach to solving ERM problems is to substitute indicators with convex upper-bounding surrogates, $\phi(z) \geq \mathbb{1}_{[z]}$ and then to minimize the resulting surrogate risk. However, such strategy generally leads to a non-convex, multi-linear optimization problem. Previous attempts to solve problems of this form have focused on computationally costly alternating optimization approaches with guarantees of convergence only to a local minimum Trapeznikov and Saligrama (2013); Wang and Saligrama (2012). Rather than attempting to solve this non-convex surrogate problem, we instead reformulate the indicator empirical risk in (4) as a maximization over sums of indicators before introducing convex surrogate. Our approach yields a globally convex upper-bounding surrogate of the empirical loss function. In the next section, we derive this reformulation for a binary tree of depth 2 in Fig. 1 before generalizing to arbitrary trees.

3.1 Simple Tree Example

Consider the decision system shown in Fig. 1. The goal is to learn the decision functions g_1, g_2, and g_3 that minimize the empirical risk (4).

In reformulating the risk, it is useful to define the "savings" for an example. The *savings*, π_k^i, for an example i, represents the difference between the worst case outcome, R_{max} and the risk $R_k(f_k, \mathbf{x}_i, y_i)$ for terminating and classifying at the kth leaf. The worst case risk is acquiring all sensors and incorrectly classifying: $R_{max} = 1 + \alpha \sum_m c_m$.

$$\pi_k^i = R_{max} - R_k(f_k, \mathbf{x}_i, y_i) = \mathbb{1}_{f_k(\mathbf{x}_i)=y_i} + \alpha \sum_{m \in S_k^C} c_m \tag{5}$$

Here, S_k^C is the complement set of sensors acquired along the path to leaf k (the sensors not acquired on the path to leaf k). Note that the savings do not depend on the decisions, $g_j's$, that we are interested in learning.

For our example, there are only 4 leaf nodes and the state of terminating in a leaf is a encoded by a product of two indicators. For instance, to terminate in Leaf 1, $g_1(\mathbf{x}_i) \leq 0$ and $g_2(\mathbf{x}_i) \leq 0$. This empirical risk can be formulated by enumerating over the leaves and their associated risks:

$$\begin{aligned} R(\mathbf{g}, \mathbf{x}_i, y_i) = & \overbrace{\left(R_{max} - \pi_1^i\right)\mathbb{1}_{g_1(\mathbf{x}_i)\leq 0}\mathbb{1}_{g_2(\mathbf{x}_i)\leq 0}}^{\text{Leaf 1}} + \overbrace{\left(R_{max} - \pi_2^i\right)\mathbb{1}_{g_1(\mathbf{x}_i)\leq 0}\mathbb{1}_{g_2(\mathbf{x}_i)>0}}^{\text{Leaf 2}} \\ & + \underbrace{\left(R_{max} - \pi_3^i\right)\mathbb{1}_{g_1(\mathbf{x}_i)>0}\mathbb{1}_{g_3(\mathbf{x}_i)\leq 0}}_{\text{Leaf 3}} + \underbrace{\left(R_{max} - \pi_4^i\right)\mathbb{1}_{g_1(\mathbf{x}_i)>0}\mathbb{1}_{g_2(\mathbf{x}_i)>0}}_{\text{Leaf 4}} \end{aligned} \tag{6}$$

Directly replacing every $\mathbb{1}_{[z]}$ with an upper bounding surrogate such as a hinge loss, $\max[0, 1+z] \geq \mathbb{1}_{[z]}$, produces a non-convex bilinear objective due the indicator product terms. Bilinear optimization is computationally intractable to solve globally.

Rather than directly substituting surrogates and solving the non-convex minimization problem, we reformulate the empirical risk with respect to the indicators in the following theorem:

Theorem 31 *The empirical risk in* (6) *is equal to* (7)*.*

$$R(g_1, g_2, g_3, \mathbf{x}_i, y_i) = R_{max} - \sum_{j=1}^{4} \pi_j^i + \max \Big[(\pi_3^i + \pi_4^i)\mathbb{1}_{g_1(\mathbf{x}_i)\leq 0} + \pi_2^i \mathbb{1}_{g_2(\mathbf{x}_i)\leq 0}, \\ (\pi_3^i + \pi_4^i)\mathbb{1}_{g_1(\mathbf{x}_i)\leq 0} + \pi_1^i \mathbb{1}_{g_2(\mathbf{x}_i)>0}, (\pi_1^i + \pi_2^i)\mathbb{1}_{g_1(\mathbf{x}_i)>0} + \pi_4^i \mathbb{1}_{g_2(\mathbf{x}_i)\leq 0}, \\ (\pi_1^i + \pi_2^i)\mathbb{1}_{g_1(\mathbf{x}_i)>0} + \pi_3^i \mathbb{1}_{g_3(\mathbf{x}_i)>0} \Big] \tag{7}$$

Proof Here, we provide a brief sketch of the proof. For full details please refer to the Appendix. We utilize the following two identities: $\mathbb{1}_{[A]}\mathbb{1}_{[B]} = \min[\mathbb{1}_{[A]}, \mathbb{1}_{[B]}]$ and $\mathbb{1}_{[A]} = 1 - \mathbb{1}_{[\bar{A}]}$ and express the risk in (6) in terms of maximizations:

$$R\left(g_1, g_2, g_3, \mathbf{x}_i, y_i\right) = R_{max} - \sum_{j=1}^{4} \pi_j^i + \pi_1^i \max\left(\mathbb{1}_{g_1(\mathbf{x}_i)>0}, \mathbb{1}_{g_2(\mathbf{x}_i)>0}\right) \\ + \pi_2^i \max\left(\mathbb{1}_{g_1(\mathbf{x}_i)>0}, \mathbb{1}_{g_2(\mathbf{x}_i)\leq 0}\right) + \pi_3^i \max\left(\mathbb{1}_{g_1(\mathbf{x}_i)\leq 0}, \mathbb{1}_{g_3(\mathbf{x}_i)>0}\right) \\ + \pi_4^i \max\left(\mathbb{1}_{g_1(\mathbf{x}_i)\leq 0}, \mathbb{1}_{g_3(\mathbf{x}_i)\leq 0}\right) \tag{8}$$

Recall that the signs of g_1, g_2, g_3 encode a unique path for $\mathbf{x}_i$. So let us consider sign patterns for each path. For instance, to reach leaf 1, $g_1 \leq 0$ and $g_2 \leq 0$. In this case, by inspection of (8), the risk is $(\pi_3^i + \pi_4^i)\mathbb{1}_{[g_1(\mathbf{x}_i)\leq 0]} + \pi_2^i \mathbb{1}_{[g_2(\mathbf{x}_i)\leq 0]} +$ constants. This is exactly the first term in the maximization in (7). We can perform such computation for each leaf (term in the max) in a similar fashion. And due to the interdependencies in (8), the term corresponding to a valid path encoding will be the maximizer in (7).

Risk Interpretability: Intuitively, in the reformulated empirical risk in (7), each term in the maximization encodes a path to one of the K leaves. The largest (active) term correspond to the path induced by the g_j's for an example $\mathbf{x}_i$. Additionally, the weights on the indicators in (7) represent the *savings lost* if the argument of the indicator is active. For example, if the decision function $g_1(\mathbf{x}_i)$ is negative, leaves 3 and 4 cannot be reached by $\mathbf{x}_i$, and therefore π_3^i and π_4^i, the savings associated with leaves 3 and 4, cannot be realized and are lost.

A distinct advantage of the reformulated risk in (7) arises when replacing indicators with convex upper-bounding surrogates of the form $\phi(z) \geq \mathbb{1}_{z\leq 0}$. Introducing such surrogates in the original risk in (6) produces a bilinear function for which a global optimum cannot be efficiently found. In contrast, introducing convex surrogate functions in (7) produces a convex upper-bound for the empirical risk.

3.2 Extension to Arbitrary Binary Trees

In this section, we generalize the empirical risk reformulation for any binary tree and present a convex surrogate. Consider a binary tree, $\mathcal{T}$, composed of $K-1$ internal nodes and K leaves. As defined in (5), each leaf has a corresponding savings π_k^i that captures the difference between the worst case risk and the risk of classifying at leaf k.

Note that in the previous example, the risk in (4) consists of a max of K terms. Each term is a weighted linear combinations of indicators, and each weight corresponds to the *savings lost* if the decision inside the indicator argument is true. For an arbitrary binary tree of K leaves, the risk has an analogous form.

Before stating the result, we define the weights for the linear combination in each term of the max. For an internal node j, we denote C_j^n as the set of leaf nodes in a subtree corresponding to a negative decision $g_j(\mathbf{x}) \leq 0$. And C_j^p is the set of leaf nodes in a subtree corresponding to a positive decision. For instance in Fig. 1, $C_1^p = \{Leaf\ 3, Leaf\ 4\}$, and in our example (7), the weight multiplying $\mathbb{1}_{[g_1(\mathbf{x}_i)\leq 0]}$ is the sum of these savings for leaves 3 and 4 (i.e. savings lost if $g_1 \leq 0$). Therefore, sets $C_j^{p/n}$ define which π_k^i's contribute to a weight for a decision term.

For a compact representation, recall that the kth rows in matrices $\mathbf{P}$ and $\mathbf{N}$ define a path to leaf k in terms of $g_1, \ldots, g_{K-1}$, and a non-zero $\mathbf{P}/\mathbf{N}_{k,j}$ indicates if $g_j \lesseqgtr 0$ is on the path to leaf k. So for each $\mathbf{x}_i$ and each leaf k, we introduce two positive weight row vectors of length $K-1$:

$$\mathbf{w}_{n,k}^i = \left[\mathbf{N}_{k,1} \sum_{l\in C_1^p} \pi_l^i, \ldots, \mathbf{N}_{k,K-1} \sum_{l\in C_{K-1}^p} \pi_l^i\right]$$

$$\mathbf{w}_{p,k}^i = \left[\mathbf{P}_{k,1} \sum_{l\in C_1^n} \pi_l^i, \ldots, \mathbf{P}_{k,K-1} \sum_{l\in C_{K-1}^n} \pi_l^i\right]$$

The jth component of $\mathbf{w}_{n,k}^i$ multiplies $\mathbb{1}_{[g_j(\mathbf{x}_i)\leq 0]}$ in the term corresponding to the kth leaf. For instance in our 4 leaf example in (7), $\left(\mathbf{w}_{n,1}^i\right)_1 = \pi_3^i + \pi_4^i$. If $\mathbf{P}/\mathbf{N}_{k,j}$ is zero then decision $g_j \gtrless 0$ is not on the path to leaf k and the weight is zero. Using these weight definitions, the empirical risk in (7) can be extended to arbitrary binary trees:

Corollary 32 *The empirical risk of tree $\mathcal{T}$ is:*

$$R(\mathbf{g}, \mathbf{x}_i, y_i) = R_{max} - \sum_{k=1}^{K} \pi_k^i$$
$$+ \max_{k\in\{1,\ldots,K\}} \mathbf{w}_{p,k}^i \begin{bmatrix} \mathbb{1}_{g_1(\mathbf{x}_i)>0} \\ \vdots \\ \mathbb{1}_{g_{K-1}(\mathbf{x}_i)>0} \end{bmatrix} + \mathbf{w}_{n,k}^i \begin{bmatrix} \mathbb{1}_{g_1(\mathbf{x}_i)\leq 0} \\ \vdots \\ \mathbb{1}_{g_{K-1}(\mathbf{x}_i)\leq 0} \end{bmatrix} \tag{9}$$

The proof of this corollary is included in the Appendix and follows the same steps as Thm. 31.

The empirical risk in (9) represents a scan over the paths to each leaf ($k = 1, \ldots, K$), and the active term in the maximization corresponds to the leaf to which an observation is assigned by the decision functions $g_1, \ldots, g_{K-1}$. An important observation is that each term in the max in (9) is a linear combination of indicators instead of a product as in (4). This transformation enables us to upper-bound each indicator function with a convex surrogate, $\phi(z)$: $\phi[g_j(\mathbf{x})] \geq \mathbb{1}_{[g_j(\mathbf{x})>0]}$, $\phi[-g_j(\mathbf{x})] \geq \mathbb{1}_{[g_j(\mathbf{x})\leq 0]}$. And the result is a novel convex upper-bound on the empirical risk in (9). We denote this risk as $R_\phi(\mathbf{g})$. And the optimization problem over a set of training examples, $\{\mathbf{x}_i, y_i\}_{i=1}^N$ and a family of decision functions $\mathcal{G}$: $\min_{\mathbf{g}\in\mathcal{G}} \sum_{i=1}^N R_\phi(\mathbf{g}, \mathbf{x}_i, y_i)$.

3.3 Linear Programming

There are many valid choices for the surrogate $\phi(z)$. However, if a hinge loss is used as an upper bound and $\mathcal{G}$ is a family of linear functions of the data then the optimization problem in (3.2) becomes a linear program (LP).

Proposition 33 *For $\phi(z) = \max(1 - z, 0)$ and linear decision functions $g_1, \ldots, g_{K-1}$, the minimization in* (3.2) *is equivalent to the following linear program:*

$$\min_{\substack{g_1,\ldots,g_{K-1},\gamma^1,\ldots,\gamma^N \\ \alpha_1^1,\ldots,\alpha_{K-1}^N,\beta_1^1,\ldots,\beta_{K-1}^N}} \sum_{i=1}^N \gamma^i \tag{10}$$

subject to:

$$\gamma^i \geq \mathbf{w}_{p,k}^i \begin{bmatrix} \alpha_1^i \\ \vdots \\ \alpha_{K-1}^i \end{bmatrix} + \mathbf{w}_{n,k}^i \begin{bmatrix} \beta_1^i \\ \vdots \\ \beta_{K-1}^i \end{bmatrix}, \quad \begin{matrix} i \in [N] \\ k \in [K], \end{matrix}$$

$$\begin{matrix} 1 + g_j(\mathbf{x}_i) \leq \alpha_j^i, \\ 1 - g_j(\mathbf{x}_i) \leq \beta_j^i, \\ \alpha_j^i \geq 0, \beta_j^i \geq 0, \end{matrix} \quad \begin{matrix} \\ i \in [N] \\ k \in [K-1] \end{matrix}.$$

We introduce the variable γ^i for each example $\mathbf{x}_i$ to convert from a maximization over leaves to a set of linear constraints. Similarly, the maximization within each hinge loss is converted to a set of linear constraints. The variables α_j^i upper-bound the indicator $\mathbb{1}_{g_j(x_i)>0}$ and the variables β_j^i upper-bound the indicator $\mathbb{1}_{g_j(x_i)\leq 0}$. Additionally, the constant terms in the risk are removed for notational simplicity, as these do not effect the solution to the linear program. For details please refer to Appendix.

Complexity: Linear programming is a relatively well-studied problem, with efficient algorithms previously developed. Specifically, for K leaves, N training points, and a maximum feature dimension of D, we have $O(KD + KN)$ variables and $O(KN)$ constraints. The state of the art primal-dual methods for LP are fast in practice, with an expected number of iterations $O(\sqrt{n}\log n)$, where n is the number of variables Anstreicher et al (1999).

Table 1 Small dataset descriptions

Name	Classes	Stage 1	Stage 2	Stage 3	Stage 4
letter	26	Pixel Count	Moments	Edge Features	-
landsat	6	Band 1	Band 2	Band 3	Band 4
SUN Mech. Turk	16	Function	Materials	Surf./Spat. Prop.	-
Image Seg.	7	Location	Pixel Int.	Color	$\cdots$

4 Experiments

We demonstrate performance of our proposed approach on two types of real world data sets. First, we demonstrate performance of the LP formulation on examples with fixed structures. In these cases, sensor subsets do not need to be learned and an exhaustive tree can be constructed over all subsets of sensors. We compare the performance of our approach to the alternating optimization scheme presented in Trapeznikov and Saligrama (2013) applied to the same tree, demonstrating efficiency and performance of our LP formulation. Next, we apply our proposed approach to data sets where dimensionality prevents exhaustive search through feature subsets and both feature subsets and tree structure must be learned along with decision functions in the tree. We show performance on classification data sets presented in Kusner et al (2014) and compare performance with Cost Sensitive Trees of Classifiers (CSTC) Xu et al (2013) and Approximately Submodular Trees of Classifiers (ASTC) Kusner et al (2014).

For all feature subset classifiers used in our proposed approach and the alternating optimization approach, classifiers are trained using logistic loss on 2^{nd}-order homogeneous polynomial expanded basis on the entire training set.

Learning Decision Functions in Fixed Structures: Fig. 2 shows performance of our proposed approach on 5 data sets where few sensors are used and fixed structure can be easily found. For the letter, landsat, and SUN Mechanical Turk data sets Frank and Asuncion (2010); Patterson and Hays (2012), a tree can be easily constructed using all possible feature subsets. For the image segmentation data set, Frank and Asuncion (2010), we fix a greedily constructed tree with 8 leaves.

For comparison, we use the alternating minimization approach proposed in Trapeznikov and Saligrama (2013). Additionally, we also show performance of a simple myopic strategy for a baseline comparison on these example. The LP approach generally performs comparably to the non-convex alternating optimization approach. Additionally, as shown in Table 2, the LP is dramatically faster during training time with comparable performance to the alternating training approach.[4] Note that we do not compare performance to ASTC or CSTC for these experiments. The purpose of these examples is to show the efficacy of the LP for training sequential decision functions independent of feature subsets, classifier design, or tree structure. Both CSTC and ASTC simultaneously learn the feature subsets, classifiers, and decision functions, limiting comparison.

Learning Tree Structure: We next applied our proposed approach to three data sets where learning a tree over all feature subsets is infeasible. For all three datasets, we learn the set of subsets as described in Section 2, with a total of

[4] All computations were performed on an Intel I5 M430 CPU @ 2.27 GHz with 4 cores.

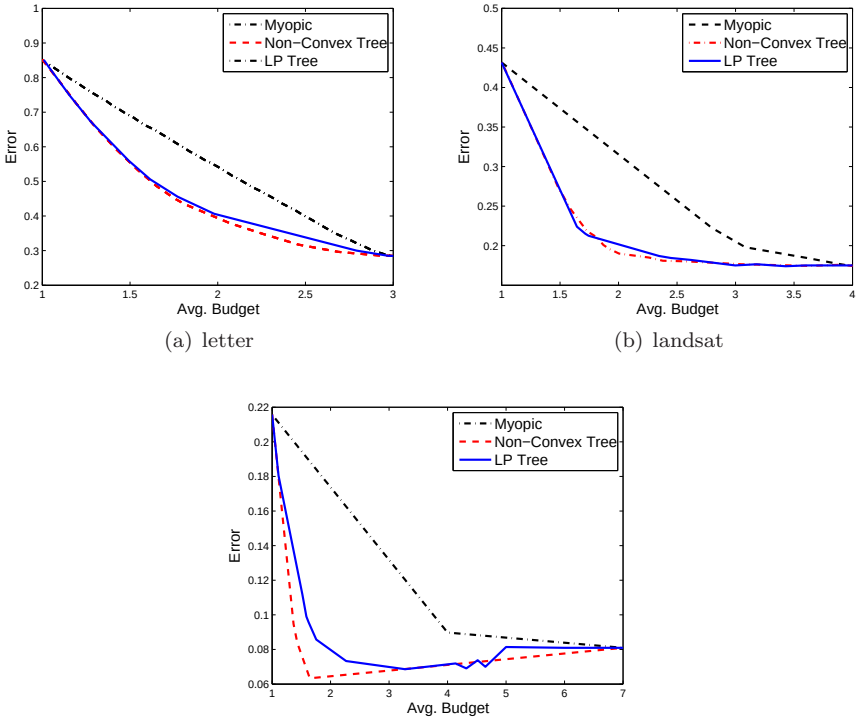

(a) letter (b) landsat

(c) Image Seg.

Fig. 2 Comparison of error vs. average budget trade-off between a myopic, AM Trapeznikov and Saligrama (2013), our LP method. LP clearly out performs the myopic approach, and generally matches or exceeds the non-convex approach with the added benefit of reduced computational cost, repeatability, and guaranteed convergence.

Table 2 Average percentage of the budget required to achieve a desired error rate chosen to be close to the error achieved using the entire set of features (approximately 95% of the improvement gained using all features compared the initial features). The percentage of the budget required is with respect to the maximum budget. The training time is the amount of time (in seconds) required to learn a policy for a fixed budget trade-off parameter α.

Dataset	Target Errors	Myopic	AM	LP	AM Train Time(sec)	LP Train Time(sec)
letter	40%	73%	48%	49%	93.56	57.03
landsat	15%	100%	75%	75%	186.0	108.7
SUN Mech. Turk	40.4%	99%	90%	90%	2818.9	71.08
Image Seg.	9%	56%	21%	26%	46.26	16.46

16 subsets of features (leaves of the tree) used. For the MiniBooNE and forest data sets, the proposed approach outperforms CSTC, with performance exceeding ASTC for the forest data set and matching ASTC for MiniBooNE. On the CIFAR dataset, the proposed LP approach matches CSTC when using 50 features, but otherwise is generally outperformed by both CSTC and ASTC. We attribute this to the limited complexity (2^{nd} order homogeneous polynomial) of the classification

functions, which requires more features to gain flexibility to accurately partition the data.

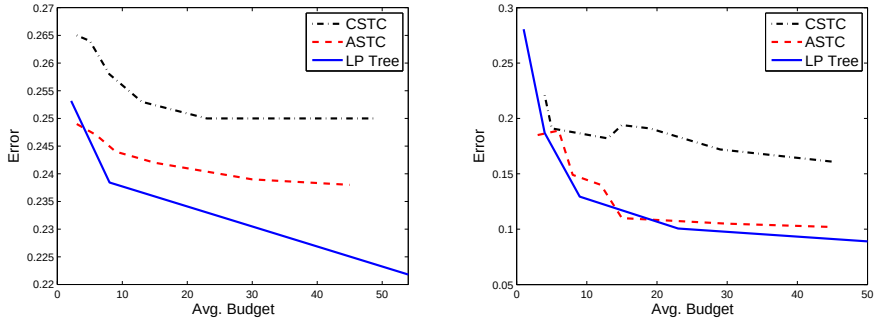

(a) Forest (b) MiniBooNE

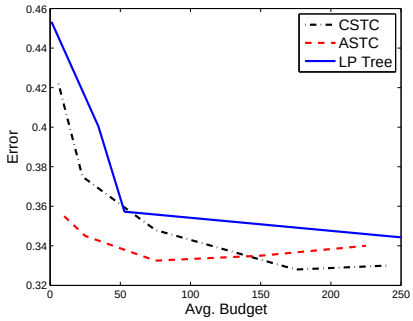

(c) CIFAR

Fig. 3 Plot of our proposed approach (LP tree) and CSTC on three real world data sets. On all three data sets, LP tree generally outperforms CSTC, producing high levels of classification accuracy with very low budgets.

5 Conclusion

A Proof of Theorem 31

The product of indicators can be expressed as a minimization over the indicators, allowing the empirical loss to be expressed:

$$
\begin{aligned}
R\left(g_1, g_{21}, g_{22}, x_i, y_i\right) = \Bigg(& R_{max} - \pi_1^i \min\left(\mathbb{1}_{g_1(x_i)\leq 0}, \mathbb{1}_{g_{21}(x_i)\leq 0}\right) \\
& - \pi_2^i \min\left(\mathbb{1}_{g_1(x_i)\leq 0}, \mathbb{1}_{g_{21}(x_i)>0}\right) - \pi_3^i \min\left(\mathbb{1}_{g_1(x_i)>0}, \mathbb{1}_{g_{22}(x_i)\leq 0}\right) \\
& - \pi_4^i \min\left(\mathbb{1}_{g_1(x_i)>0}, \mathbb{1}_{g_{22}(x_i)>0}\right) \Bigg).
\end{aligned}
$$

By swapping the inequalities in the arguments of the indicator functions, the minimization functions can be converted to maximization functions:

$$\begin{aligned} R\left(g_1, g_{21}, g_{22}, x_i, y_i\right) = R_{max} &+ \pi_1^i \max\left(\mathbb{1}_{g_1(x_i)>0}, \mathbb{1}_{g_{21}(x_i)>0}\right) - \pi_1^i \\ &+ \pi_2^i \max\left(\mathbb{1}_{g_1(x_i)>0}, \mathbb{1}_{g_{21}(x_i)\leq 0}\right) - \pi_2^i + \pi_3^i \max\left(\mathbb{1}_{g_1(x_i)\leq 0}, \mathbb{1}_{g_{22}(x_i)>0}\right) \\ &- \pi_3^i + \pi_4^i \max\left(\mathbb{1}_{g_1(x_i)\leq 0}, \mathbb{1}_{g_{22}(x_i)\leq 0}\right) - \pi_4^i \Bigg). \end{aligned}$$

Note that due to the dependence of the indicators, there will always be 3 maximization terms equal to 1 and 1 maximization term equal to zero. As a result, the sum of maximizations can be expressed as a maximization over the 4 possible combinations, yielding the expression:

$$\begin{aligned} R\left(g_1, g_{21}, g_{22}, x_i, y_i\right) = \Bigg(& R_{max} - \pi_1^i - \pi_2^i - \pi_3^i - \pi_4^i \\ & + \max\Big((\pi_3^i + \pi_4^i)\mathbb{1}_{g_1(x_i)\leq 0} + \pi_2^i \mathbb{1}_{g_{21}(x_i)\leq 0}, (\pi_3^i + \pi_4^i)\mathbb{1}_{g_1(x_i)\leq 0} + \pi_1^i \mathbb{1}_{g_{21}(x_i)>0}, \\ & (\pi_1^i + \pi_2^i)\mathbb{1}_{g_1(x_i)>0} + \pi_4^i \mathbb{1}_{g_{21}(x_i)\leq 0}, (\pi_1^i + \pi_2^i)\mathbb{1}_{g_1(x_i)>0} + \pi_3^i \mathbb{1}_{g_{21}(x_i)>0}\Big)\Bigg). \end{aligned}$$

B Proof of Corollary 32

The product of indicators over an arbitrary binary tree is given by:

$$R(\mathbf{g}, \mathbf{x}_i, y_i) = \sum_{k=1}^{K} \overbrace{R_k(f_k, \mathbf{x}_i, y_i)}^{\text{risk of leaf } k} \underbrace{\prod_{j=1}^{K-1} [\mathbb{1}_{g_j(\mathbf{x}_i)>0}]^{\mathbf{P}_{k,j}} [\mathbb{1}_{g_j(\mathbf{x}_i)\leq 0}]^{\mathbf{N}_{k,j}}}_{\text{state of } G_k(\cdot) = \mathbf{x}_i \text{ in a tree}}.$$

Converting the product into a minimization over indicators, the function can be rewritten:

$$R(\mathbf{g}, \mathbf{x}_i, y_i) = \sum_{k=1}^{K} \left(R_{max} - \pi_k^i\right) \min_{j\in\{1,\ldots,K-1\}} \left([\mathbb{1}_{g_j(\mathbf{x}_i)>0}]^{\mathbf{P}_{k,j}}, [\mathbb{1}_{g_j(\mathbf{x}_i)\leq 0}]^{\mathbf{N}_{k,j}}\right)$$

and using the identity $\mathbb{1}_A = 1 - \mathbb{1}_{\bar{A}}$, this can be converted to the maximization:

$$R(\mathbf{g}, \mathbf{x}_i, y_i) = R_{max} - \sum_{k=1}^{K} \pi_k^i + \sum_{k=1}^{K} \pi_k^i \max_{j\in\{1,\ldots,K-1\}} \left([\mathbb{1}_{g_j(\mathbf{x}_i)\leq 0}]^{\mathbf{P}_{k,j}}, [\mathbb{1}_{g_j(\mathbf{x}_i)>0}]^{\mathbf{N}_{k,j}}\right).$$

As in the 2-region case, the dependence of the indicators always results in $K-1$ maximization terms equal to 1 and 1 maximization term equal to 0. By examination, the sum of maximization functions can be expressed as a single maximization over the paths of the leaves, resulting in a loss shown in (9).

C Additional Explanation of Prop. 4.1

The linear program of Prop. 33 is constructed by replacing the indicators with hinge-losses of the appropriate signs:

$$\min_{\substack{g_1,\dots,g_{K-1},\gamma^1,\dots,\gamma^N \\ \alpha_1^1,\dots,\alpha_{K-1}^N,\beta_1^1,\dots,\beta_{K-1}^N}} \sum_{i=1}^{N} \gamma^i \tag{11}$$

subject to:

$$\gamma^i \geq \mathbf{w}_{p,k}^i \begin{bmatrix} \alpha_1^i \\ \vdots \\ \alpha_{K-1}^i \end{bmatrix} + \mathbf{w}_{n,k}^i \begin{bmatrix} \beta_1^i \\ \vdots \\ \beta_{K-1}^i \end{bmatrix}, \quad \begin{matrix} i \in [N] \\ k \in [K] \end{matrix},$$

$$\begin{matrix} 1 + g_j(\mathbf{x}_i) \leq \alpha_j^i, \\ 1 - g_j(\mathbf{x}_i) \leq \beta_j^i, \\ \alpha_j^i \geq 0, \beta_j^i \geq 0, \end{matrix} \quad \begin{matrix} \\ i \in [N] \\ k \in [K-1] \end{matrix}.$$

Note that the linear program arises based on the fact that any maximization can be converted to a linear constraint with the introduction of a new variable. The maximization in the objective for each observation is replaced by the introduction of the variable γ^i and the first constraint. The maximization functions in the hinge losses are replaced by the remaining constraints, introducing the variables $\alpha_j^i = \max(1 + g_j(\mathbf{x}_i), 0)$ and $\beta_j^i = \max(1 - g_j(\mathbf{x}_i), 0)$, respectively.

References

Anstreicher K, Ji J, Potra F, Ye Y (1999) Probabilistic analysis of an infeasible-interior-point algorithm for linear programming. Math Oper Res 24(1):176–192

Awasthi P, Feldman V, Kanade V (2013) Learning using local membership queries. In: Conference on Learning Theory

Bilgic M, Getoor L (2007) Voila: Efficient feature-value acquisition for classification. In: AAAI, vol 22, p 1225

Busa-Fekete R, Benbouzid D, Kegl B (2012) Fast classification using sparse decision dags. In: ICML

Chakaravarthy V, Pandit V, Roy S, Awasthi P, Mohania M (2007) Decision trees for entity identification: ap- proximation algorithms and hardness results. In: PODS

Chen M, Xu Z, Weinberger KQ, Chapelle O, Kedem D (2012) Classifier cascade: Tradeoff between accuracy and feature evaluation cost. In: AISTATS, pp 235–242

Cicalese F, Laber E, Saettler A (2014) Diagnosis determination: decision trees optimizing simultaneously worst and expected testing cost. In: ICML, pp 2124–2132

Dulac-Arnold G, Denoyer L, Preux P, Gallinari P (2011) Datum-wise classification: a sequential approach to sparsity. In: MLKDD, Springer, pp 375–390

Efron B, Hastie T, Johnstone I, Tibshirani R (2004) Least angle regression. The Annals of statistics 32(2):407–499

Frank A, Asuncion A (2010) UCI machine learning repository. URL `http://archive.ics.uci.edu/ml`

Gao T, Koller D (2011) Active classification based on value of classifier. In: NIPS, vol 24, pp 1062–1070

He H, Daume III H, Eisner J (2012) Imitation learning by coaching. In: Advances In Neural Information Processing Systems (NIPS), pp 3158–3166

Ji S, Carin L (2007) Cost-sensitive feature acquisition and classification. Pattern Recognition 40(5):1474–1485

Kanani P, Melville P (2008) Prediction-time Active Feature-Value Acquisition for Cost-Effective Customer Targeting. In: Advances In Neural Information Processing Systems (NIPS)

Kapoor A, Horvitz E (2009) Breaking Boundaries: Active Information Acquisition Across Learning and Diagnosis. In: Advances In Neural Information Processing Systems (NIPS)
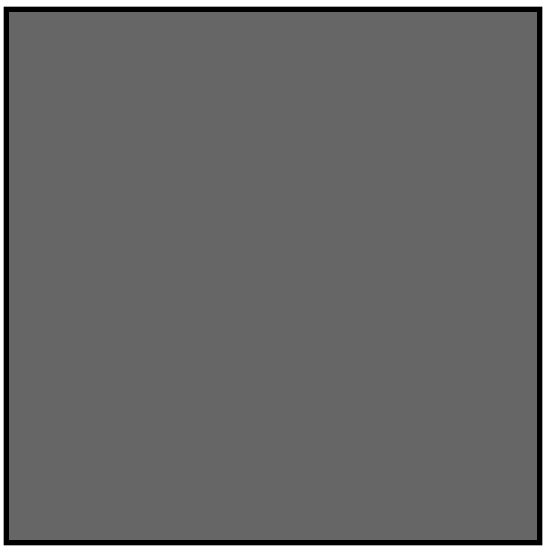
Karayev S, Fritz M, Darrell T (2013) Dynamic feature selection for classification on a budget. In: ICML Workshop on Prediction with Sequential Models
Kusner M, Chen W, Zhou Q, Xu E, Weinberger K, Chen Y (2014) Feature-cost sensitive learning with submodular trees of classifiers
Patterson G, Hays J (2012) Sun attribute database: Discovering, annotating, and recognizing scene attributes. In: CVPR
Sheng VS, Ling CX (2006) Feature value acquisition in testing: A sequential batch test algorithm. In: International Conference on Machine Learning (ICML), pp 809–816
Trapeznikov K, Saligrama V (2013) Supervised sequential classification under budget constraints. In: AISTATS
Trapeznikov K, Saligrama V, Castañon DA (2013) Multi-Stage Classifier Design. In: Machine Learning, pp 1–24
Wang J, Saligrama V (2012) Local supervised learning through space partitioning. In: Advances in Neural Information Processing Systems 25
Wang J, Saligrama V (2013) Locally-linear learning machines (l3m). In: Asian Conference on Machine Learning, pp 451–466
Wang J, Bolukbasi T, Trapeznikov K, Saligrama V (2014a) Model selection by linear programming. In: European Conference on Computer Vision, pp 647–662
Wang J, Trapeznikov K, Saligrama V (2014b) An lp for sequential learning under budgets. In: AISTATS
Xu Z, Chapelle O, Weinberger K (2012) The greedy miser: Learning under test-time budgets. In: ICML
Xu Z, Kusner M, Chen M, Weinberger K (2013) Cost-sensitive tree of classifiers. In: ICML-13
Zhang C, Zhang Z (2010) A Survey of Recent Advances in Face Detection. Tech. rep., Microsoft Research
Zubek VB, Dietterich TG (2002) Pruning Improves Heuristic Search for Cost-Sensitive Learning. In: International Conference on Machine Learning (ICML), pp 19–26